Trinocular Correspondence for Particles and Streaks

J. K. Kearney

Department of Computer Science
The University of Iowa

Abstract

This paper examines the problem of matching corresponding object points in three views of a scene. The matching problem is a critical step in the recovery of 3-dimensional position by stereo image processing and has important application in the analysis of particle distributions and fluid flows. We introduce an algorithm that iteratively searches for all possible pairwise matches using trinocular consistency constraints. The algorithm is shown to be equivalent to a search for edges that belong to all perfect matchings in a bipartite graph that links consistent matches between two views. Three such problems are solved in parallel. An attractive property of the matching algorithm is its graceful degradation in response to image distortion and modeling error. Although measurement errors may prevent successful matching, wrong matches can almost always be avoided if the error in the image position of a particle can be bounded. Thus, noise can cause a loss of acuity but should not cause the introduction of gross misinformation that could result from incorrect matching. Experiments with synthetic particle images demonstrate that this approach can lead to significantly greater numbers of matches than algorithms that match in a single direction.

This work was supported in part by National Science Foundation Grant IRI-8808896 and Office of Naval Research Grant ONR 00014-88K-0632.

1. Introduction

Accurate measurement of particle location and motion is important for a diverse assortment of scientific and engineering problems. A particle is a small, compact object such as an oil droplet, a blood cell, or a piece of dust. Fluid motions can be quantitatively assessed by tracking particles seeded in flow fields. The three-dimensional distribution of particles at an instant in time can also give important information about the characteristics of gas and fluid motions. In this paper we focus on the problem of measuring the location of large numbers of particles at a snapshot in time using image processing techniques. An extension of the method to measure particle motions is discussed.

Information about the three-dimensional location of an object can be recovered from multiple images of the object taken simultaneously from different viewpoints.[1] This process is usually performed in three steps. First, sets of object features are independently found in each image. An object feature is a visually distinctive characteristic that can be reliably detected from different viewpoints. The edges and corners of surface markings are often used as features. Next, features are matched across the images. The last step is to estimate the three-dimensional location of the feature by triangulation. The use of two or more views to infer three-dimensional structure and shape is called stereo vision. Considerable research has been directed towards methods that rely on two cameras. Recently, a number of papers have pointed out the advantages of introducing a third camera.[2-11]

This paper addresses the matching problem for particles. The particles to be viewed are assumed to be small spheres with a size near the resolution limit of the camera. For our model, we will assume that particles appear as points in an image. Our problem is to find points in two or more different images that are projections of the same particle. This presents a difficult matching problem, because all image points look alike. Further, neighboring image points may lie at disparate depths. As a consequence, stereo matching algorithms that rely on surface continuity or the appearance of distinctive markings are inappropriate for particle matching.

We present a general algorithm for matching image points in three views. The method relies solely on geometric constraints based on the camera arrangement. In previous formulations of the trinocular matching algorithm, one image was used to verify the consistency of possible matches between points in the other two images. Our approach generalizes this strategy by considering all pairwise matches between points in all images. The matching problem is shown to be equivalent to a search for edges contained in all perfect matchings of a bipartite graph. The bipartite graph links points between two images that satisfy a geometric consistency test. To take maximal advantage of the geometric constraints, three such problems must be solved. Our algorithm employs an iterative search to find certain matches. The algorithm was tested with synthetically generated particle images. The results show that it finds significantly greater numbers of matches than earlier methods when matching uncertainty is high.

An attractive property of the algorithm is its graceful degradation in response to image distortion and modeling error. Although measurement errors may prevent successful matching, wrong matches can almost always be avoided if the error in the image position of a particle can be bounded. Thus, noise can cause a loss of acuity but should not cause the introduction of gross misinformation that could result from incorrect matching. The density of matches obtained by the method can be optimized by regulating the number of particles seeded in the volume and by selectively arranging the cameras to minimize ambiguity in the matching process. The method can be simply extended to match image streaks formed by using long exposure views of moving particles.

2. Three Camera Correspondence

Before describing point matching among three views, we consider stereo matching of points between two views. The stereo correspondence problem is facilitated by geometric constraints when the relative positions and orientations of the cameras are known. Given the position of a point in one image, its location in a second image is known to lie on a line. As pictured in Figure 1, a three-dimensional particle P that projects to point p_r in the right image is known to lie on the ray originating at the camera center, c_r, and passing through the image at p_r. The epipolar line is defined to be the locus of points on the left image to which the ray $\overline{c_r p_r}$ projects. Given the position and orientation of the left camera with respect to the right, we can confine our search for correspondences to the epipolar line.[12] The epipolar line is the intersection of the image plane with a plane defined by the two camera centers and the image point p_r. This plane is called the epipolar plane. Every particle that lies in the epipolar plane defined by (c_r, c_l, p_r) will appear as a point on the epipolar line in the left image and could be consistently matched to p_r. If there is a single particle in the epipolar plane, then there will be a single consistent point in the left image and it can be unambiguously matched to p_r. When two or more particles lie on an epipolar plane, additional constraints must be brought to bear on the matching process.

A third camera provides additional information that can be used to discriminate among the consistent match points.[2-11] Figure 2 illustrates the mutual constraints among three images. Consider the pair of points (p_{l_i}, p_{c_i}) selected from the left and center images, respectively. Assume that both points lie on the epipolar line defined by the other so that they could consistently be matched. If p_{l_i} and p_{c_i} correspond to the same particle, then their common match in the right image must lie at the intersection of two epipolar lines. The two points define independent epipolar constraints in the right image. If there is no point at the intersection, then

p_{l_j} can be eliminated from consideration as a candidate match for p_{c_i}. When there is a point p_{r_k} at the intersection of the epipolar line in the right image, then we say the set of points $(p_{l_j},p_{c_i}p_{r_k})$ is a **consistent triplet**. We define a consistent triplet as a set of three points, one from each image, such that each point lies on the epipolar line defined by the other two points.

The appearance of a point at the intersection of the epipolar lines in the right image is a necessary but not a sufficient condition to infer correspondence. A potential ambiguity arises when three particles lie along three image rays that intersect at a point. Figure 3 demonstrates how three image points might be falsely interpreted as projections of a particle located at the intersection of the rays. However, when this occurs there must be more than one consistent triplet that includes p_{l_i}. (We assume all points are visible in all views.) If there is a single triplet of points $(p_{l_i},p_{c_j},p_{r_k})$ containing p_{l_i} among the set of all triplets satisfying the three camera epipolar constraints, then we can be certain that the three points correspond to the same particle. We call such a triplet a **matched triplet**. Whenever a matched triplet is found, all other consistent triplets containing p_{c_j} and p_{r_k} can be eliminated from consideration. This may enable new matched triplets to be declared. By iteratively searching for matched triplets and removing impossible contenders from the set of consistent triplets, the number of matched triplets may successively grow. This process can continue until no new matched triplets are found on an iteration.

We illustrate the utility of excluding impossible triplets using the particle distribution pictured in Figure 3. The three particles appear as points p_{l_1}, p_{l_2}, and p_{l_3} in the left image. In the center image the three points are labeled p_{c_1}, p_{c_2}, and p_{c_3}. And in the right image they are called p_{r_1}, p_{r_2}, and p_{r_3}. From the three camera epipolar constraints the set of consistent triplets is restricted to

a $(p_{l_1},p_{c_1},p_{r_1})$

b $(p_{l_2},p_{c_2},p_{r_2})$

c $(p_{l_3},p_{c_3},p_{r_3})$

d $(p_{l_1},p_{c_2},p_{r_3})$.

Because *b* is the only triplet among the consistent triplets that includes p_{l_2}, we know that it is a matched triplet. This allows us to eliminate *d* from the set of possible triplets, because p_{c_2} has been associated with a matched triplet and hence, cannot correspond to any other point. This leaves only *a* with p_{l_1} as a member and only *c* with p_{l_3}. Consequently, we can identify *a*, *b*, and *c* as matched triplets.

The selection of matched triplets from among the set of consistent triplets relies on a uniqueness assumption. We assume that every image point has exactly one correct match in every other image. If there is only one consistent triplet containing a point, then this triplet must contain the correct matches for that point. Thus, a consistent triplet can be determined to be a matched triplet if any of the three points appears in no other consistent triplet. Once the triplet is confirmed as a set of matching points, no member of the triplet can be matched to any other point. Therefore, all other triplets containing the matched points can be deleted from the set of consistent triplets.

For the problem at hand, it is not even necessary to uniquely determine all three points in a match triplet. The 3-dimensional location of a particle can be estimated knowing its projection in only two images. Hence, we can recover a particle's location if we can determine with certainty any two members of its match triplet. It is still desirable to find all three members whenever possible, because our estimate of particle position will likely be improved with the redundant information provided by a third point. However, the set of consistent triplets may only provide enough information to determine two of the three correspondents. We illustrate this situation with the following set of consistent triplets:

$$\begin{array}{ll} a & (p_{l_1}, p_{c_1}, p_{r_1}) \\ b & (p_{l_1}, p_{c_1}, p_{r_2}) \\ c & (p_{l_2}, p_{c_2}, p_{r_1}) \\ d & (p_{l_2}, p_{c_2}, p_{r_2}) \end{array}$$

Only consistent triplets a and b contain p_{l_1}. In both of these triplets p_{l_1} is matched with p_{c_1}. This pairwise match can be determined even though the correspondent in the right image cannot be discerned. Combinations of consistent triplets similar to this arise when image locations are inexactly known. If the epipolar constraint is modified to tolerate image measurement errors, then one arrangement that could lead to the set above is illustrated in Figure 4. In the next section we discuss adaptations necessary to accommodate imaging errors.

Our matching strategy is to search for pairs of matching points. For each pair of images, we collapse the set of consistent triplets into the set of all possible point pairs by dropping the entry from the third image out of every triplet. There are three pairwise combinations of three images, so we have three matching problems to solve. For the moment, we concentrate on solving each of these problems separately.

The set of possible matches between points in two images corresponds to a bipartite graph, where the vertices of the graph are image points. A vertex belongs to one of two disjoint sets based on the image in which the point resides. Edges connect points across the two sets. Two points are connected if there is at least one consistent triplet containing both points.

A perfect matching in a bipartite graph is defined to be a subset of the edges such that every vertex is connected to exactly one other vertex.[13] The set of edges that link corresponding points in the two images defines one perfect matching. There must be exactly one such edge for each point because we assume thai all particles are visible in all images. Unfortunately, there may be many possible perfect matchings. We need to find edges that belong to all possible perfect matchings. We can state with certainty that points connected by these edges correspond to the projections of the same particle in 3-dimensional space.

Given a bipartite graph that contains a perfect matching, any vertex with a single edge incident on it must belong to all perfect matchings. For our problem, the two vertices connected by the edge represent points in two images that correspond to the same particle. Once a correct match is determined, it can be recorded and both vertices and the edge that connects them can be deleted from the graph. When a vertex is deleted, all other edges incident on this vertex can also be deleted. None of these edges, called forbidden edges, can belong to any perfect match. When a forbidden edge is deleted in this fashion, the number of edges incident on some vertex remaining in the graph is reduced by one. If there remains only one edge incident on the vertex, then it too can be matched. In this way, correct matches can be successively found and removed from the graph. The process ends when no more vertices of degree one are left in the bipartite graph.

There exist more general methods to eliminate forbidden edges from a bipartite graph.[13] These techniques are able to find forbidden edges not found by our method. Deletion of these edges may, in turn, lead to the discovery of more correct matches. The more general process of detecting forbidden edges may lead to exponentially complex searching and hence is impractical when large numbers of particles are used. We will be content to detect forbidden edges only when they are incident on matched vertices.

To this point, we've treated the three bipartite graphs as independent. However, the edges are derived from a single set of consistent triplets and whenever a edge between two points is forbidden, then all triplets containing the pair of vertices connected by the edge can be eliminated. Thus, deletion of an edge in one bipartite graph permits at least one edge in each of the three bipartite graphs to be deleted.

3. Ambiguity and Error

The camera models required to calculate the epipolar lines must be estimated and consequently contain error. Image discretization and sensor distortion also contribute noise to the measurements. Because of this error, we must extend the the epipolar line to be a narrow band and accept all points within this band as consistent with the epipolar constraint.

An important property of the approach outlined above is its resilience to distortion of image positions. If the error in image location can be bounded, then deviations should not contribute to the mismatching. Matching errors can only occur when

(a) the correct triplet is not among the set of possible matches

and

(b) there is at least one set of three image points consistent with the geometric constraints that are not projections of the same particle.

There are two conditions which may cause the correct match to be absent from the set of candidate matches. The correct point may not be visible or its location may be so severely perturbed by error that it violates the epipolar constraint. A particle may not appear on the image because low contrast and noise hamper discriminability, it is out of the field of view, or it is occluded by another particle. With proper lighting and camera positioning the first two situations can be avoided. If two particles happen to lie on a ray passing through a camera center, the particle nearer the camera will hide the particle to the rear. The single image point will consistently match points corresponding to the occluded and occluding particles in the other two images. Since it actually represents two physical points, either of these choices would be correct. If this dual point is matched to one of its two correspondents, it will be eliminated from consideration for the second. This will leave a point with no correct match and it could potentially be coupled with a geometrically consistent, but invalid point. This problem could be avoided by strengthening the matching criteria. If we require that both of the points to be matched have the other as the sole candidate among all consistent triplets, then the dual point will not be matched and it will always provide contention against incorrect matches. However, strengthening the matching criteria can lead to substantial reduction in the total number of matches. We will assume that all points are visible in all views, recognizing that the rare event of particle occlusion could possibly lead to a matching error.

The distortion of the image location of particles is incorporated into the epipolar test to guarantee that the correct match is always acceptable. If the position distortion can be bounded, then the algorithm can be adjusted so that the correct match is guaranteed to satisfy the epipolar constraints. Consequently, the correct match will always provide competition against incorrect matches.

Clustering and coincidental alignment of particles can always create ambiguity in the matching process. Matching errors will not be made, but information may be lost due to the inability to discriminate between plausible candidates. The problem is worsened as the magnitude of errors to be accommodated grows. Let the displacement of a particle from its epipolar line be bounded by ε. The single error term, ε incorporates distortion from the nominal image position caused by optical and discretization error and error in the epipolar line due to inaccuracies in the camera models. To assure that the correct point is considered, any point lying within ε of the epipolar line must be accepted as a possible match. All points in a 3-dimensional wedge will project to image locations within ε of the epipolar line. The wedge of confusion can be obtained by rotating the epipolar plane about an axis located at the image center and parallel to the epipolar line. The rotation must extend ε above and below the epipolar line on the image.

For a triplet to consistently satisfy the broadened acceptance criteria when accommodating error, each point must lie within ε of the epipolar lines defined by the two other members of the triplet. The region of acceptance on the image is a trapezoid centered on the intersection of the epipolar lines. The area of 3-dimensional space that projects within the admissible trapezoid image is defined by the intersection of two wedges. A confound exists whenever a point lies within the skewed pyramid formed by intersecting the two wedges. Figure 4 illustrates an arrangement of two particles that leads to an unresolvable ambiguity in the rightmost image. Either point in the right image could be consistently associated with points p_{l_1} and p_{c_1} in the left and center images. Either point could also be associated with points p_{l_2} and p_{c_2}. There is insufficient information to match p_{r_1} or p_{r_2} to points in either the left or right images. The likelihood of an ambiguity increases with the density of points and the size of acceptance region. The volume of the wedge intersection is determined by the error tolerance and the angle between the epipolar lines. More nearly parallel epipolar lines will lead to larger regions of intersection.

Careful arrangement of the cameras can contribute to favorable conditions for matching. The two epipolar lines in the right image will be parallel whenever the point to be matched lies on the plane defined by the three camera centers. In this situation the third camera provides no additional constraint. If the cameras are arranged so that the plane of the camera centers lies outside the sampled volume, then the epipolar lines can never be parallel.

4. The Three-Camera Matching Algorithm

The first step in the matching algorithm is to determine all consistent triplets. Recall that a triplet of image points is consistent if each point lies sufficiently close to the epipolar lines defined by the other two points. The criteria for nearness to the epipolar line must bound errors due to modeling and measurement errors.

The second step is to create and initialize the data structures used in the iterative search for matches. We assign an index i = 1,2 or 3 to each image. Consistent triplets will be ordered by the image indices. Thus, the triplet (x,y,z) contains p_{1_x} from image 1, p_{2_y} from image 2, and p_{3_z} from image 3. We define a circular ordering on the three images, so that each image has a left neighboring image and a right neighboring image. The functions $left(i)$ and $right(i)$ return the indices of the left and right neighboring images. For a point p_{i_j} in image i, we want to find a match $p_{left(i)_k}$ in the image to the left and a match $p_{right(i)_l}$ in the image to the right. We keep a record of the status of these matches as the algorithm progresses. Initially, a point is labeled as unmatched in both of its neighboring images. Whenever a match is discovered, the two matching points are labeled matched in the appropriate images.

We also construct three adjacency lists for the hypergraph whose vertices are the points in the three images and whose edges are the consistent triples. Each list corresponds to a vertex-edge list. For each vertex j in image i, we associate all ordered triples that contain the value j in the i^{th} position.

The last step is to iteratively search the adjacency lists for each image looking for certain matches in the other two images. When all consistent triplets on the adjacency for point p_{i_j} contain the same point $p_{left(i)_k}$ in position $left(i)$, then p_{i_j} is marked as matched in image $left(i)$ and $p_{left(i)_k}$ is marked as matched in image i. Forbidden edges incident on $p_{left(i)_k}$ are eliminated by traversing its adjacency list and deleting all consistent triplets with an entry other than j in the i^{th} position. Because every consistent triplet is represented on the adjacency lists for all its members, deletion of the triplet requires removal from all three adjacency lists. Finally, the adjacency list for p_{i_j} is examined for possible matches in image $right(i)$. As with the left image, whenever a match is found, the matched points are marked and forbidden edges are deleted. The complete algorithm is presented in Figure 5.

```
(1)  Determine all consistent triplets;

(2)  Mark all points as unmatched in all images;
     Create adjacency lists for all points in images 1,2, and 3;
(3)  Repeat
     For each image $i = 1,2,3$
          For each point $p_{i_j}$ in image $i$
               If $p_{i_j}$ is unmatched in image $left(i)$ and there is a single candidate $p_{left(i)_k}$ in the
                  $left(i)^{th}$ position for all consistent triplets on $p_{i_j}$'s adjacency list
                    Then [
                         mark $p_{i_j}$ as matched in image $left(i)$;
                         mark $p_{left(i)_k}$ as matched in image i;
                         delete forbidden edges;
                         ]
               If $p_{i_j}$ is unmatched in image $right(i)$ and there is a single candidate $p_{right(i)_k}$ in
                  the $right(i)^{th}$ position for all consistent triplets on $p_{i_j}$'s adjacency list
                    Then [
                         mark $p_{i_j}$ as matched in image $right(i)$;
                         mark $p_{right(i)_k}$ as matched in image i:
                         delete forbidden edges;
                         ]
          EndFor
     EndFor
   Until no matches new found on one iteration
```

Figure 5. The three-camera matching procedure

5. Testing The Method

Three matching algorithms were tested with synthetically generated images. The algorithm presented in Figure 5 was compared with two versions of an approach that uses the third image only to confirm the consistency of matches between the two principal images. We will refer to this as biased matching because points in all images are not treated with equal status. The more complex form of this algorithm is presented in Figure 6. The simpler form of the algorithm performs only one iteration of the repeat loop. Thus, it considers each point in image 1 only once.

```
Repeat
    For each unmatched point $p_{1_i}$ in the center image
        For each unmatched point $p_{2_j}$ near $p_{1_i}$'s epipolar line in the left image
            If there is at least one point in image 3 satisfying the epipolar constraints for
               $p_{1_i}$ and $p_{2_j}$
               Then let $p_{2_j}$ be a candidate match set for $p_{1_i}$
        EndFor
        If $p_{1_i}$ has a single candidate match $p_{2_j}$
            Then mark $p_{1_i}$ and $p_{2_j}$ as matched
    EndFor
Until no points matched on this iteration
```

Figure 6. The iterative version of the biased three-camera matching procedure.

The use of synthetic images enabled us to examine the performance of the iterative matching scheme with varied amounts of controlled error. Synthetic models of particle distributions were created by randomly selecting points in a unit volume. The modeled particles were projected onto three image planes. Cameras were modeled as ideal pinhole devices. The camera focal lengths could be set to any value, but were assumed to be the same for the three cameras. The positions of the cameras relative to the observation space could be adjusted by the experimenter. Camera gaze was automatically determined so that the optic axis passed through the center of the volume. The image coordinates of points in the world were calculated using perspective projection. A sample model and the corresponding images are graphically displayed in Figure 7.

Error was introduced into the image position data by surrounding image points with a circle of uncertainty. The matching algorithm assumed that the true image value could be located anywhere within a circle of radius ε. All sources of error were lumped into the single error term, ε. Any point within ε of an epipolar line was considered to satisfy the epipolar constraint.

For the experiments reported here, we positioned the cameras along the axes of the world coordinate system located at the center of the volume bounding particle positions. Cameras were placed at a distance of 2 from the origin in units normalized for the dimension of the observation space. Camera gaze was automatically directed toward the origin. Consequently, each image plane was parallel to a side of the bounding volume. The focal lengths of the cameras were all set to .05. This created images roughly equivalent to that obtained with a 35mm camera observing a one meter volume positioned to fill the field of view.

The matching algorithms were tested with a simulated volume containing 100 particles. Values of ε were chosen to approximate image resolutions between $\frac{1}{25}$ and $\frac{1}{500}$.

The results for the three matching techniques are displayed in Figure 8. The number of correct, *unambiguous matches determined* by each method is graphed in Figure 8.a. Each data point plotted on the graph was averaged over 5 tests. The simulation experiments highlighted the robustness of the three camera matching strategy. No matching errors were made by any of the methods in any experiment.

The results demonstrate the influence of uncertainty on matching. With small levels of error, all of the methods found matches for all points. As the error increased, fewer matches were determined. At high levels of error, few or no matches were unambiguously found. The pattern of decline in number of matches with increasing error was similar for all three methods. However, the unbiased matching technique tolerated significantly higher levels of error than either of the biased methods. The iterative biased method led to a minor improvement over the one-pass biased method.

The average execution times for the matching portions of the three algorithms (step 3 in Figure 5), are graphed in Figure 8.b. The matching programs were implemented in Mathematica on a Macintosh IIx. The graph shows that the improved performance of the unbiased method comes with significant cost in efficiency.

6. Three Camera Streak Matching

The matching algorithm can be simply extended to match streaks resulting from particle motion over a long exposure view of a moving particle. Our approach is to approximate the streak by a line segment connecting the endpoints of the streak. Streak endpoints can be determined by hand or by using a contour estimation algorithm. For reasonably short time intervals the line segment will be a good approximation to particle velocity. A consistent triplet of steaks consists of a coupled pair of consistent point triplets. The three camera epipolar constraints apply to both endpoints. Thus, a streak triplet is consistent if both endpoints of every streak satisfy the epipolar constraints for the respective endpoints of the other two streaks.† A match is declared whenever there is a single consistent match in another image.

As with point matching, mismatching can be avoided if the distortion of steak endpoints is bounded. As a fortunate byproduct of matching streaks, the dual constraint provided by the two endpoints reduces ambiguity in the matching process. A confound exists only when anòther streak has endpoints that are confounded with *both* endpoints of the streak to be matched.

7. Summary

Many of the tasks for which particle distribution and motion are important are conducted as laboratory experiments. The three camera matching technique is well-suited for these applications. Cameras can be arranged to optimize discriminability and particle density can be controlled to maximize the number of matches. The technique offers the potential to locate and track large numbers of points in complex flows. The matching procedure is simple and remarkably robust.

References

1. J.R. Dhond and J.K. Aggarwal, "Structure from stereo - A review," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 19, no. 6, pp. 1489-1510, November 1989.
2. V. J. Milenkovic and T. Kanade, "Trinocular vision using photometric and edge orientation constraints," *Proceedings of the Image Understanding Workshop*, p. 163, December 1985.

†A streak's endpoint in one image may correspond to either endpoint of a streak from another image. Consequently, two endpoint orderings must be considered for each streak. If the a streak satisfies the epipolar constraints in both of its orderings, then we will form two consistent streak triplets.

3. A. Gerhard, H. Platzer, J. Steurer, and R. Lenz, "Depth extraction by stereo triples and a fast correspondence estimation algorithm," *8th International Conference on Patte;n Recognition*, p. 512, October 1986.

4. M. Ito and A. Ishii, "Range and shape measurement using three-view stereo analysis," *Proceedings Computer Vision and Pattern Recognition*, p. 9, 1986.

5. M. Pietikainen and D. Harwood, "Depth from three camera stereo," *Proceedings Computer Vision and Pattern Recognition*, p. 2, 1986.

6. M. Yachida, U. Kitamura, and M. Kimachi, "Trinocular vision: New approach for correspondence problem," *8th International Conference on Pattern Recognition*, p. 1041, October 1986.

7. M. Yachida, "3-D data acquisition by multiple views," *Robotics Research: Third International Symposium*, pp. 11-18, MIT Press, Cambridge, MA, 1986.

8. E. Gurewitz, I. Dinstein, and B. Sarusi, "More on the benefit of a third eye for machine stereo perception," *Proceedings of the Eighth International Conference on Pattern Recognition*, pp. 966-968, October 1986.

9. Y. Ohta, T. Yamaoto, and K. Ikeda, "Collinear trinocular stereo using two-level dynamic programming," *Proceedings of the Ninth International Conference on Pattern Recognition*, pp. 658-662, November 1988.

10. S. Yokoi, G. Medioni, and R. Nevatia, "Segment-based trinocular vision," *Proceedings of the USA-JAPAN Symposium on Flexible Automation - Crossing Bridges: Advances in Flexible Automation and Robotics*, pp. 817-821.

11. Y. Kitamura, K. Masatoshi, and M. Yachida, "Trinocular vision using edge continuity and local features," *Systems and Computers in Japan*, vol. 19, no. 12, pp. 12-23, December 1988.

12. B.K.P Horn, *Robot Vision*, The MIT Press, Cambridge, MA, 1986.

13. L. Lovasz and M.D. Plummer, in *Matching Theory*, Elsevier Science Publishers, Amsterdam, The Netherlands, 1986.

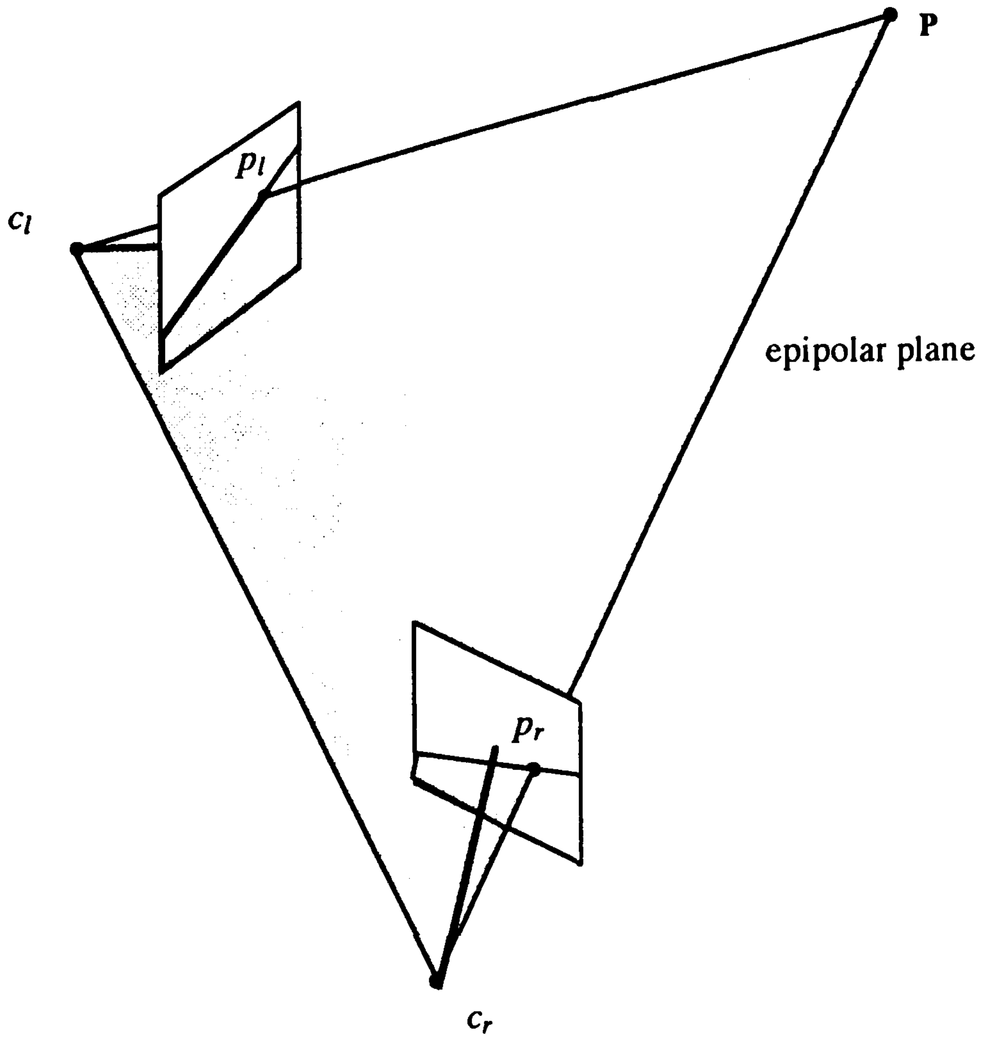

Figure 1. The epipolar constraint for two images.

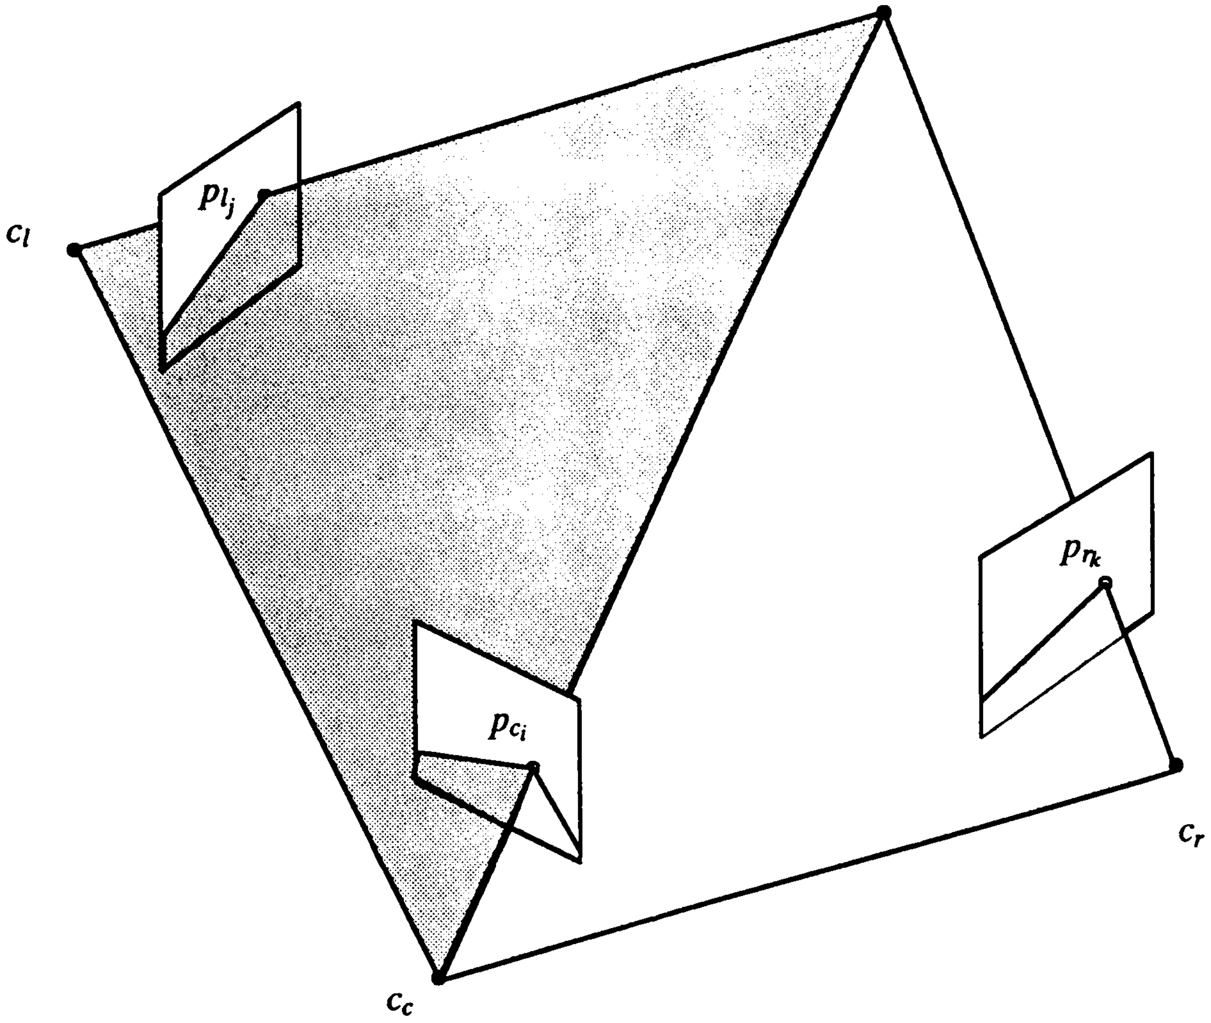

Figure 2. Mutual epipolar constraints among three images.

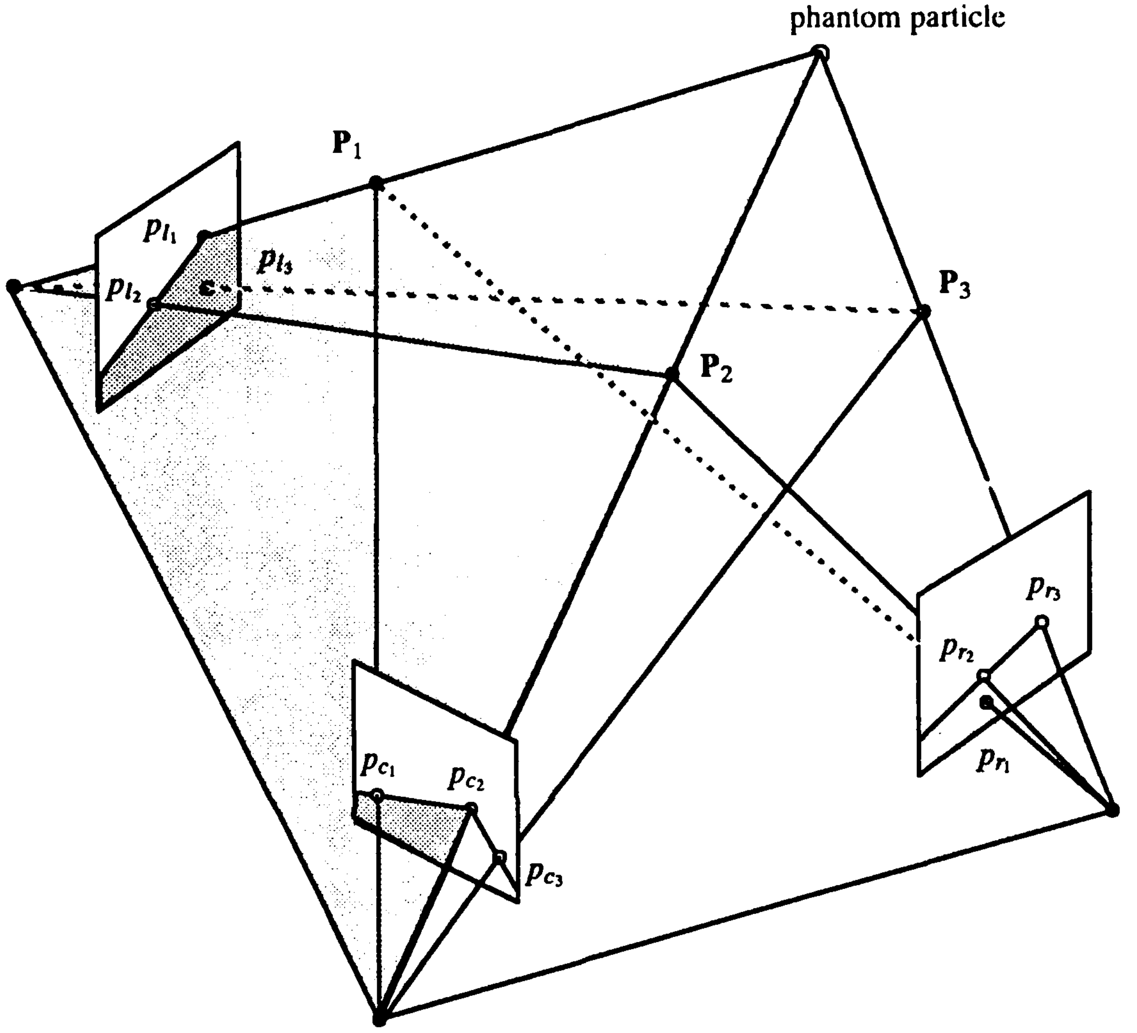

Figure 3. The three image points $(p_{l_1}, p_{c_2}, p_{r_3})$ coincidentally satisfy the 3-camera constraints for the phantom particle. Points should only be matched when there is a consistent triplet containing a point that is a member of no other consistent triplet among the set of possible matches.

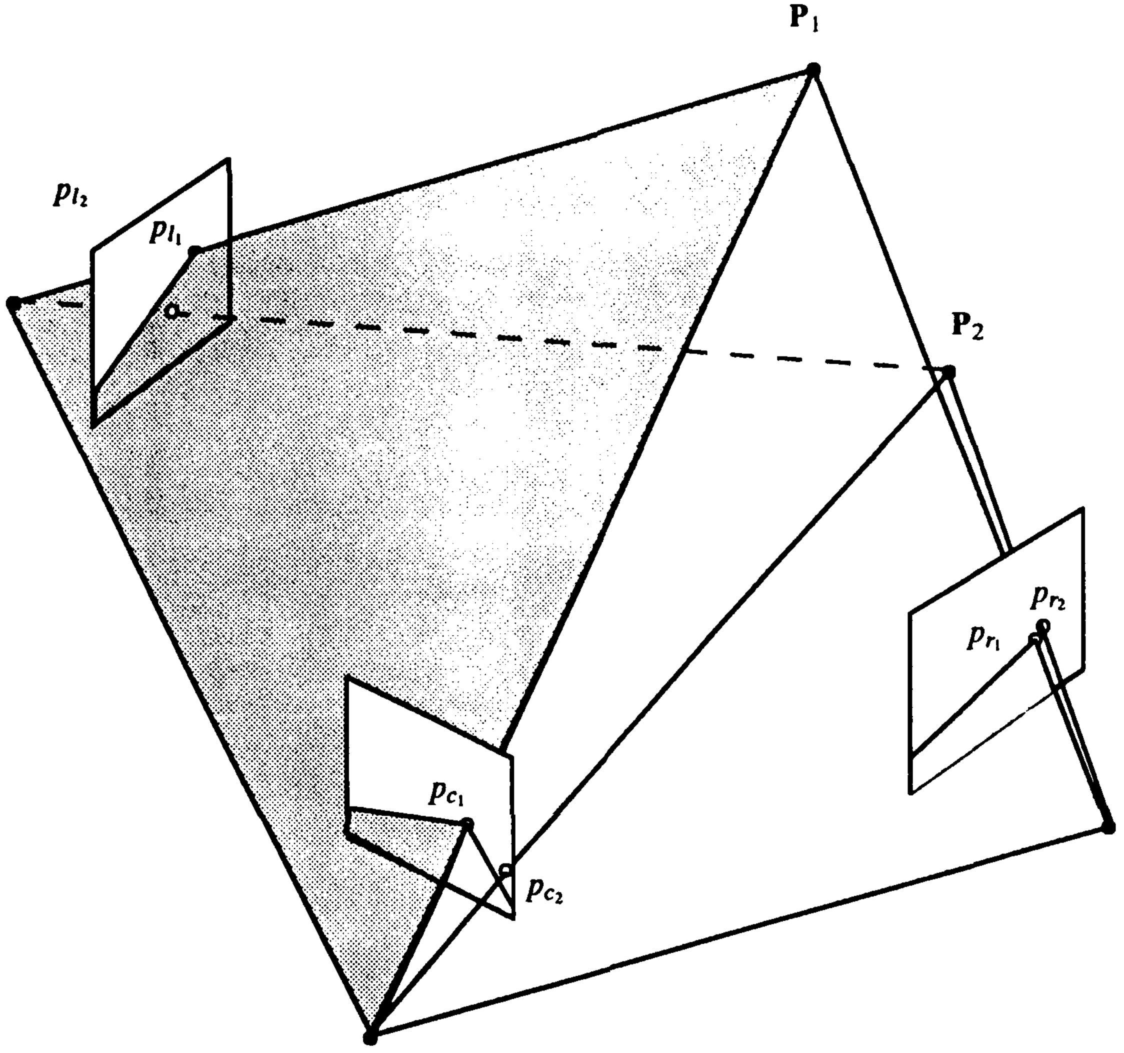

Figure 4. A configuration in which points can be matched only between the left and center images. Because of measurement errors, the correspondent in the right image cannot be determined.

THIS

PAGE

IS

MISSING

IN

ORIGINAL

DOCUMENT

Figure 5
And 6

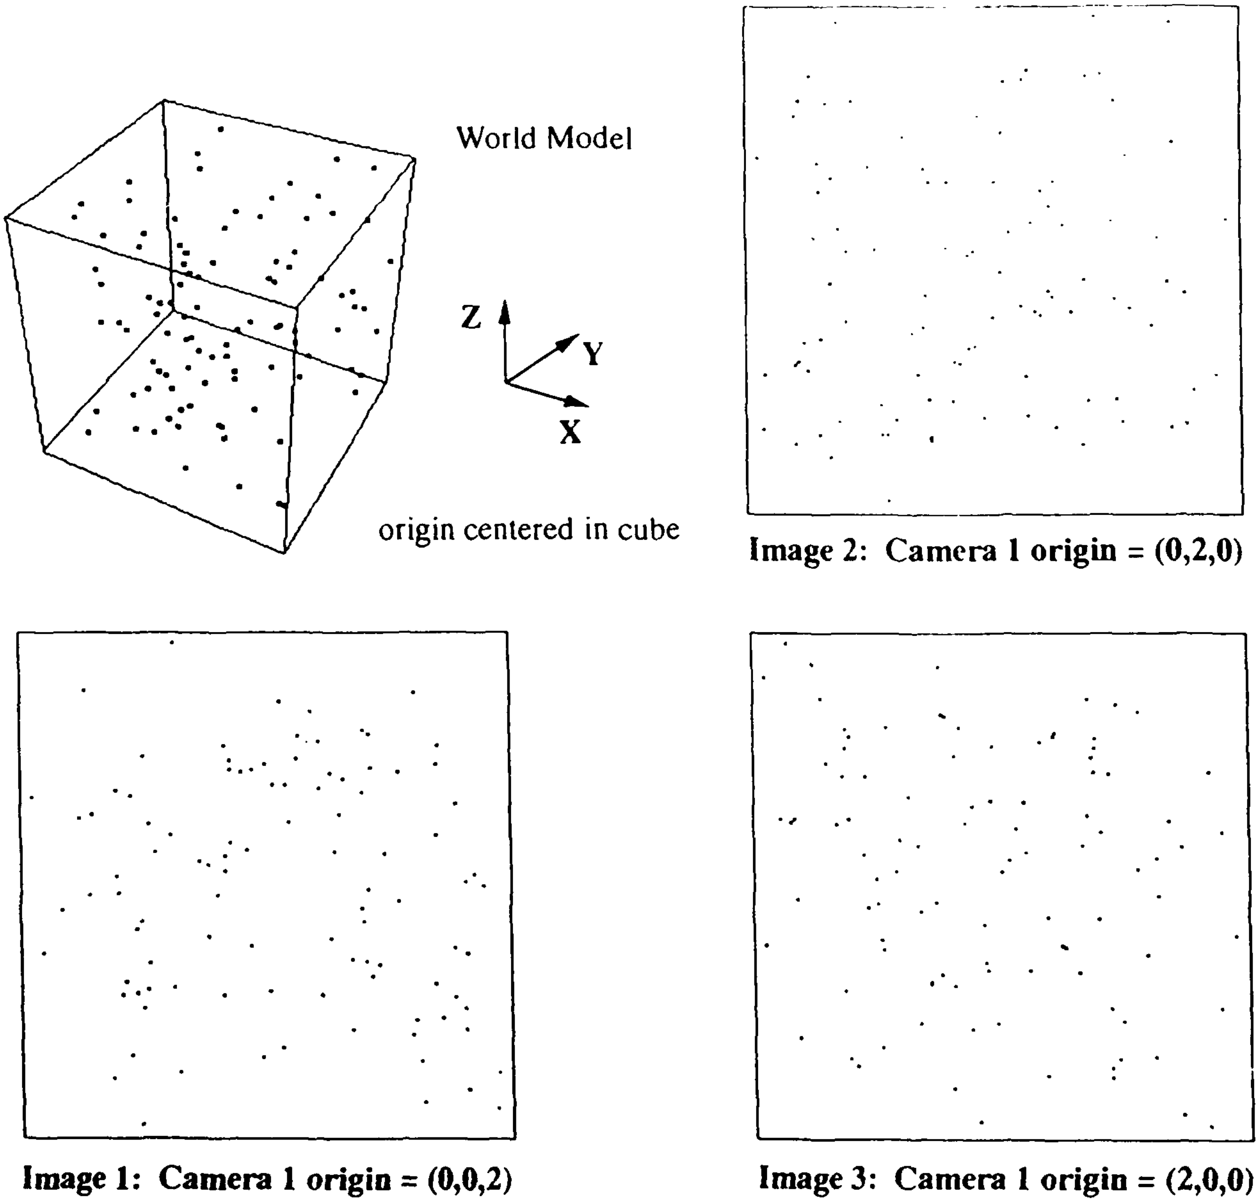

Image 2: Camera 1 origin = (0,2,0)

Image 1: Camera 1 origin = (0,0,2)

Image 3: Camera 1 origin = (2,0,0)

Figure 7. A simulated model of 100 particles and the corresponding images from three orthogonal views.

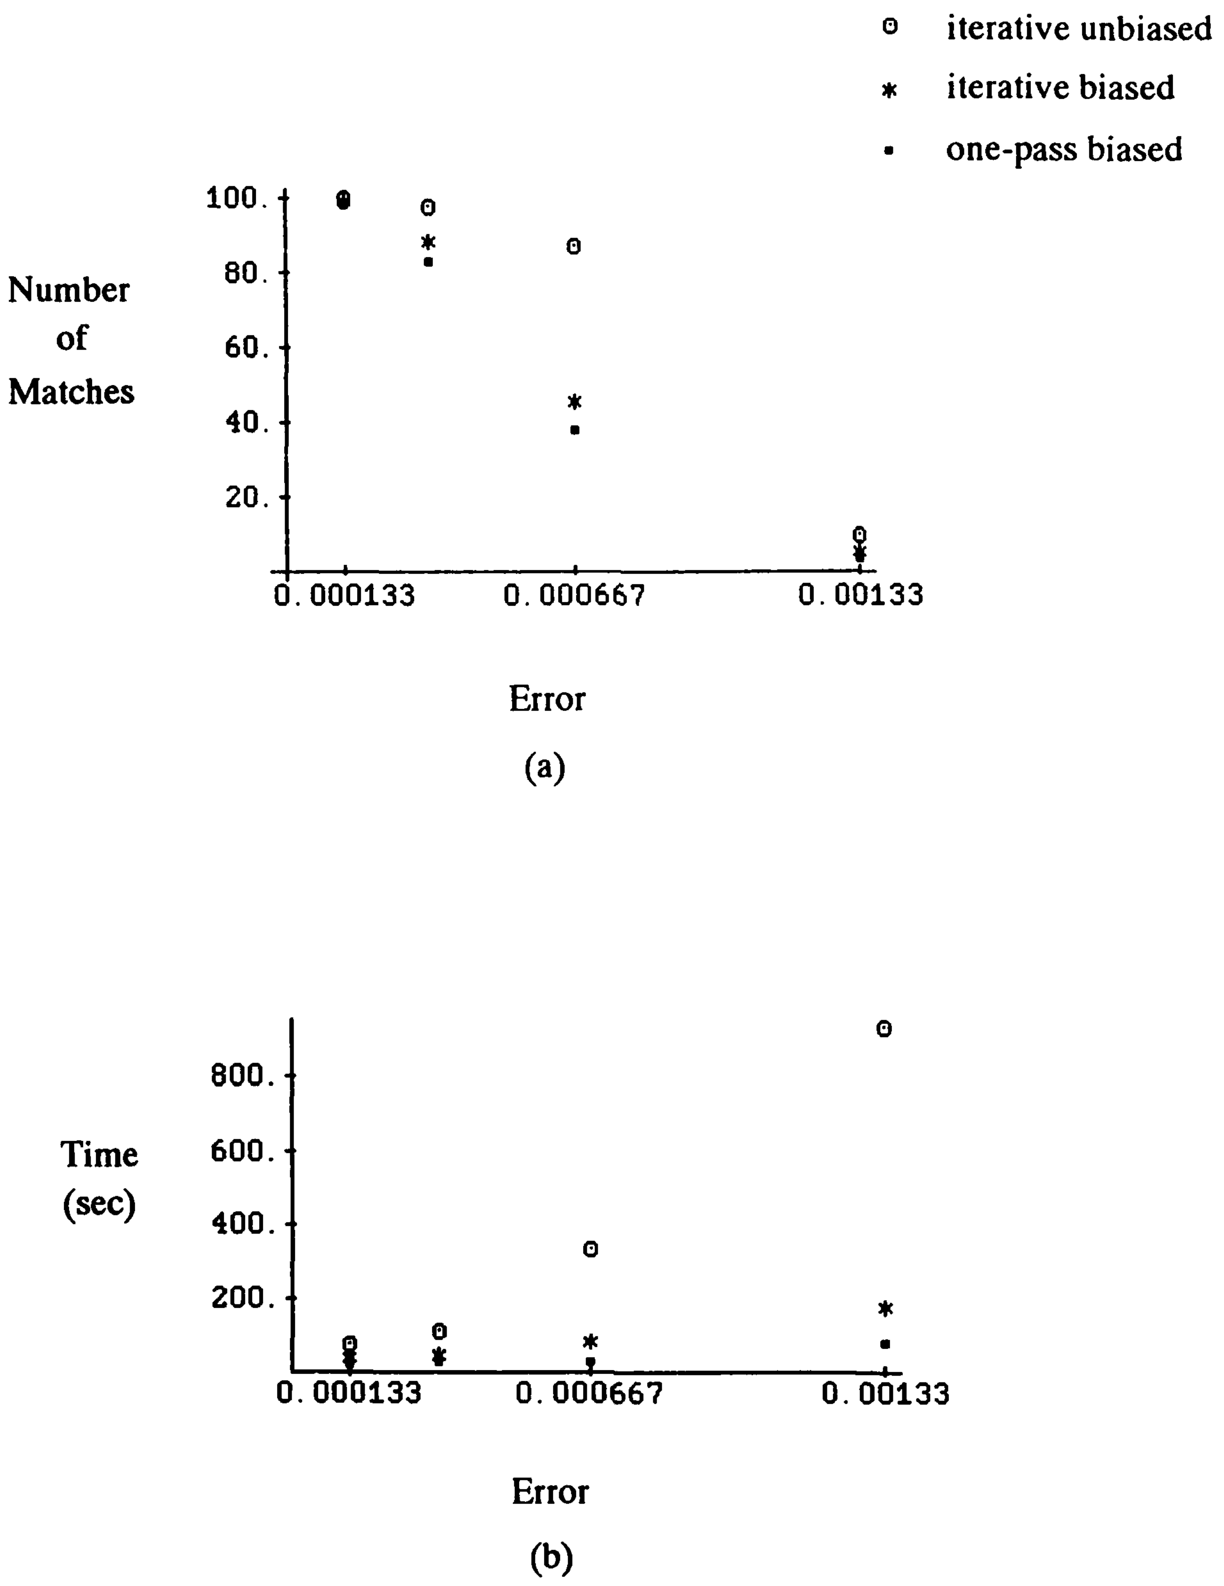

(a)

(b)

Figure 8 The results of experiments with synthetic particle images.